Model Checking Probabilistic ATL with Imperfect Recall and Imperfect Information

Bryan Kumara

2 September 2019

1 Abstract

We find the complexity of PATL with regards to imperfect information and imperfect recall in an explicitly represented system. We begin by motivating our logic by an example of the board game Risk, showing that it is a multi-agent probabilistic game. We add a restriction of being able to only see adjacent tiles as a case for imperfect information. Following that, we fix the syntax and semantics of our language and return to show how this fits with our motivating example.
Model-checking PATL with imperfect information and imperfect recall can be done by way of a linear algorithm. We first sort out the different cases. The more complex cases rely on guessing a strategy σ_A, then pruning our model, and finally solving the resulting PCTL formula. Hence, the complexity of PATL with imperfect information and imperfect recall is Δ_2^p.

2 Introduction

Consider a modified version of the board game Risk. In this game, we have two to six players played over a map of Earth separated into 42 territories and 6 continents [6]. Players take turn in playing their troops and to conquer adjacent territories determined by dice rolls. Players are also allowed to form coalitions to act towards a specific goal, with the ultimate goal of occupying every territory on the board and eliminate other players. We currently ignore the card mechanics of Risk, and focus solely on infantry troops attacking and moving based on dice rolls. Cases like this are what Probabilistic Alternating-Time Temporal Logic (PATL) model. They are probabilistic in nature as the player's movement outcomes are determined by dice rolls. Players are also able to form coalitions and their movements are prescribed by the map (adjacent tiles only). Risk is an example of a game of imperfect recall, as how we get to step S_i depends only on S_{i-1}. Traditionally, in Risk we can see how the whole board is set up, which is an example of perfect information [7]. So, let us modify Risk such that players are only given information of tiles adjacent to theirs. This creates a game of imperfect information as players do not know the conditions of all states at times during the game. To describe this situation, we can augment PATL with imperfect information and imperfect recall [3]. In this paper, we investigate the model checking complexity to our logical system[4].

2.1 Related Works

We discuss works that are relevant to PATL and its model checking techniques. [1] provided a foundation to construct our language while [2] and [3] helped construct our definition of imperfect information and imperfect recall leading to our system. All three of the paper also explored model checking techniques for probabilistic systems. For more in-depth information regarding model checking of probabilistic systems, we refer the reader to [5] and [6].

3 Probabilistic Alternating-time Temporal Logic

In this section, we introduce PATL with imperfect information and imperfect recall ($PATL_{ir}$) along with its models, Probabilistic Concurrent Game Structure with Incomplete Information ($PCGS_i$). We assume a familiarity with Measure Theory to explain some of the theory concerning Markov Decision Processes.

3.1 Syntax

The syntax of $PATL_{ir}$ is given below:

Definition 1 (PATL)

$$\varphi ::= p \mid \neg\phi \mid \psi \wedge \phi \mid \langle\langle A\rangle\rangle \left[X\phi\right]_{\bowtie r} \mid \langle\langle A\rangle\rangle \left[\phi U\psi\right]_{\bowtie r} \mid \langle\langle A\rangle\rangle \left[\phi R\psi\right]_{\bowtie r}$$

where $p \in AP$, $A \subseteq Ag$, $\bowtie \in \{\leq, <, =, >, \geq\}$, and $r \in [0,1]$

Intuitively, we let $\langle\langle A\rangle\rangle \left[\psi\right]_{\bowtie r}$ express that players of a coalition A are able to collectively enforce ψ with a probability $\bowtie r$. We can define the other basic logical connectives $\vee, \rightarrow, \leftrightarrow$ in terms of $\wedge$ and $\neg$. We let $X\phi$ to express that ϕ is true on the next state, $\phi U\psi$ to mean that ϕ holds at least until ψ becomes true, and $\phi R\psi$ to mean that ϕ holds only until ψ does. Similarly, we define other operators like $F\psi$ in terms of our 3 fundamental operators.

3.2 Semantics

We let the set of probability distributions on X is denoted by $\mathcal{D}(X)$. Our model is defined as follows:

Definition 2 *A $PCGS_i$ is a tuple $\mathcal{G} = \langle Ag, S, AP, \lambda, Act, d, PI, \delta, \sim_a\rangle$ such that:*

- *Ag is the set of agents such that agent $a \in Ag$, and a coalition A is $A \subseteq Ag$.*

- *S is the set of states such that state $s \in S$.*
- *AP is the set of atomic propositions such that p and $q \in AP$.*
- *$\lambda : S \to \wp(AP)$ is the* labelling function.
- *Act is the set of actions.*
- *d is a protocol function $d : (Ag \times S) \to (\wp(Act_a) \setminus \emptyset)$ which indicates a set of actions available to a at a specific state. We require that if $s \sim_a s'$ then $d(a,s) = d(a,s')$, that is, agents have the same actions available in indistinguishable states.*
- *PI be the initial distribution $PI : S \to [0,1]$ such that $\sum_{s \in S} PI(s) = 1$.*
- *δ is a probabilisitic transition function such that for each state s and each joint action $\langle j_1, \ldots, j_k \rangle$, a probabilistic transition function δ returns the conditional probability $\delta(s'|s, \langle j_1, ..., j_k \rangle)$ of a transition from s to s' if each player $a \in Ag$ chooses move j_a. A joint action $\langle j_1, \ldots, j_k \rangle$ at s is such that $j_a \in d(a,s)$ for each agent a. Given a state s, $D(s)$ is the set of joint actions available in s.*
- *for every $a \in Ag$, the* indistinguishability relation *$\sim_a$ is an equivalence relation on S.*

A path of $\mathcal{G}$ from state s is an infinite sequence of states $\pi = s, s_0, s_1, ...$ and we let Ω be the set of all paths, with Ω_s be the set of all paths starting from s.

Having defined our semantics, we now define the concept of a strategy and outcome with regards to imperfect information and imperfect recall.

We define imperfection information by the indistinguishability relation $\sim_a \subseteq S \times S$, one per agent $a \in Ag$. If we have $s \sim_a s'$ for states s and s' then agent a cannot distinguish between s and s'. $\sim_a$ is an equivalence relation. Assume a $PCGS_i$ $\mathcal{G}$. An imperfect information strategy (ir-strategy) is:

Definition 3 (Uniform Memoryless Strategy) *A (uniform memoryless) strategy for player a is a function $\sigma_a : S \mapsto C(a)$, for $C(a) = \bigcup_{s \in S} d(a,s)$ the set of all actions available for player a at any state s, such that for all states s, s':*

1. *$\sigma_a(s) \in d(a,s)$ (the strategy has to return available actions);*
2. *if $s \sim_a s'$ then $\sigma_a(s) = \sigma_a(s')$ (the strategy has to return the same action in indistinguishable states).*

The set of an agent a's ir-strategy is represented by Σ_a^{ir} and we write Coalition A's collective ir-strategy as $\Sigma_A^{ir} = \prod_{a \in A} \Sigma_a^{ir}$. A strategy σ_a prescribes only moves available to a. We let Σ_a denote the set of all strategies available to a. If we have σ_A, we write the set of possible outcomes of σ_A starting at a state s as $Outcomes(s, \sigma_A)$ of measures that agents in A enforce when σ_A is followed and the probability distribution $O_s^{\sigma_A}$ ranges over $Outcomes(s, \sigma_A)$.

Let a coalition A and $\overline{A}$ have strategies σ_A and $\sigma_{\overline{A}}$. Once a starting state s and strategies have been chosen, the game is reduced to a Markov Decision Process. The probability of a path can be represented by a transition probability function.

We can also define an observation function from $\sim_a$:

Definition 4 (Observation Function)

An observation function $O_a : S \to O$ such that for every agent a and state s, $O_a(s)$ is some observation $o \in O$. Then $s \sim_a s'$ iff $O_a(s) = O_a(s')$ so we also have $O_a(s) = \{s' \mid s \sim_a s'\}$.

Let us return to the Markov Decision Process. Often, we will want to know what the probability is of reaching a set of events of interest to us and so we introduce reachability probabilities:

Definition 5 (Reachability Probability) *Let $\Diamond B$ represent a set of events that are of interest. Then:*

- *$\Diamond B$ is a measurable set of paths and agrees with the union of all basic cylinders $Cyl(s_0...s_n)$.*
- *$s_0...s_n$ is a path fragment in an MDP $\mathcal{M}$ such that $s_0...s_{n-1} \notin \Diamond B$ and $s_n \in \Diamond B$. $paths_{fin}(\mathcal{M}) \cap (S \setminus B) * B$ to represent the set of all paths fulfilling that condition and the cylinder sets are disjoint.*
- *The probability of reaching B: $Pr^{\mathcal{M}}(\Diamond B) =$ $\Sigma_{s_0...s_n \in paths_{fin}(\mathcal{M}) \cap (S \setminus B) * B} Pr^{\mathcal{M}}(Cyl(s_0...s_n)) =$ $\Sigma_{s_0...s_n \in paths_{fin}(\mathcal{M}) \cap (S \setminus B) * B} \iota_{init} s_0 \cdot P(s_0...s_n)$.*
- *The probability of reaching a B state in n steps in $\mathcal{M}$ through states C only is given by: $Pr^{\mathcal{M}}(C \cup^{\leq n} B) = \Sigma_{t \in B} \Theta_n^{M'}(t)$.*

We let $M' = M_{B \cup (S \setminus (C \cup B))}$ and $\Theta_0^{\mathcal{M}} = \iota_{init}$. For in-depth understanding, we refer the reader to [5].

Having defined the necessary concepts, we now define the semantics formally. We write $\mathcal{G}, q \models_{ir} \varphi$ to show that state q satisfies φ in the structure $\mathcal{G}$ with imperfect information and imperfect recall. We omit $\mathcal{G}$ when it is clear we are referring to $\mathcal{G}$.

Definition 6 *The satisfaction relation $\models_{ir}$ is defined for all states and paths of $\mathcal{G}$ inductively as follows:*

$q \models_{ir} p$ $\quad$ $p \in \lambda(q)$
$q \models_{ir} \neg\phi$ $\quad$ *iff* $q \not\models_{ir} \phi$
$q \models_{ir} \phi \wedge \psi$ $\quad$ *iff* $q \models_{ir} \phi$ *and* $q \models_{ir} \psi$
$q \models_{ir} \langle\langle A \rangle\rangle [X\psi]_{\bowtie r}$ *iff for some ir-strategy* $\sigma_A \in \Sigma_A^{ir}$,
$\quad$ $O_s^{\sigma_A}(\pi \in \Omega_s \mid \pi[2] \models_{ir} \psi) \bowtie r$
$\pi \models_{ir} \langle\langle A \rangle\rangle [\phi U \psi]_{\bowtie r}$ *iff for some ir-strategy* $\sigma_A \in \Sigma_A^{ir}$,

$$
\begin{aligned}
& O_s^{\sigma_A}(\pi \in \Omega_s \mid \textit{for some } i \geq 0,\ \pi[i] \models_{ir} \phi \textit{ and for any} \\
& 0 \leq j < i, \pi[j] \models_{ir} \phi) \bowtie r. \\
\pi \models_{ir} \langle\langle A \rangle\rangle\, [\phi R \psi]_{\bowtie r} \ & \textit{iff for some ir-strategy } \sigma_A \in \Sigma_A^{ir}, \\
& O_s^{\sigma_A}(\pi \in \Omega_s \mid \textit{for some } i \geq 0,\ \pi[i] \models_{ir} \phi \textit{ and for all} \\
& 0 \leq j \leq i,\ \pi[i] \models_{ir} \psi, \textit{ or for all } 0 \leq j \leq i,\ \pi[i] \models_{ir} \psi) \bowtie r.
\end{aligned}
$$

3.3 $PCGS_i$ and Risk

Having formally defined our logical system, we now return to illustrate that this does fit our motivating example. The following picture illustrates Risk with ir:

Figure 1: Risk board game as seen by the red player

Before we proceed, an explanation about the colours are needed. We assume the role of the red player, who is currently picking the Middle East to attack adjacent regions. This gives out 4 arrows to show which territory to attack. The territory available for attack from the Middle East is coloured bright green while other territories are more muted. The muted red colours mean that the Red player is unable to make any attacks. For the territory in Mexico, it is unable to attack Latin America as there is only one troop stationed on that territory.

We see that in a P-iCGS $\mathcal{G} = \langle Ag, S, AP, \lambda, Act, d, PI, \delta, \sim_a \rangle$ for the above:

- There can be 3 agents in Ag, the red, yellow, and green player.
- The set of states consist of consists of functions $f : T \to (Ag \times \mathbb{N})$ for each territory which states each player a has $x \in \mathbb{N}$ troops in that territory. For instance: the red player has 4 troops in Greenland.
- AP consist of statements such as "Red player wins" and "All of Green player's troops in Australia are destroyed".

- *Act* consists of functions $attack(T, m)$ where T is a territory and m is the number of troops. For instance: "Red player attacks South Africa territory with 2 troops".

- The protocol function limits available actions. For instance, Red cannot attack Great Britain from Greenland because they are not adjacent territory. We can formalise it as such: For agent a and state s, the protocol function is $d(a, s) = \{attack(T, m) \mid \text{where s controls some } T' \text{ adjacent to } T\}$.

- Instead of equivalence relation, an observation function is equivalent and provides a better illustration. An observation function in Risk is best demonstrated as Red noticing that West Africa has 14 troops stationed: $O_a(s) = \{T' \mid a \text{ controls some } T \text{ adjacent to } T'\}$,

- PI is our initial distribution, where players have their starting territories and troops. For a game of 3 agents we have $\binom{42}{14}\binom{28}{14}(14^{105})$ equal probabilities, or around $(4,076667183)(10^{149})$ equally possible initial distributions.

- δ comes when a player invades a territory, and the outcome is determined by dice rolls until the territory is either defended or gained. The attacker may play 3 dices and the defender 2, depending on the number of troops available. The highest outcomes determines which player loses their units. For instance, Red is attacking Green from Middle East to India. The probability that Red wins the territory is 57.87%. There does not seem to be a general formula to compute the probabilities of Risk, hence a table from Data Genetics serves to illustrate the probabilities of different attackers and defenders. For δ, only one player is making a move, the rest are just defending. So, in a function $\delta(s'|s, \langle j_a, ..., j_n\rangle)$, if player a is making a move $j_a = attack(T, m)$, then $\forall i \neq a, j_n = \text{idle}$. Note that at the end of every turn, players receive armies based on territories they control, with extra amounts handed to them for each continent they control. For instance, a player controlling Australia gains 2 additional infantry per turn.

Figure 1 can also help illustrate the semantics of $PATL_{ir}$:

1. Let us illustrate $\langle\langle A\rangle\rangle\, [X\psi]_{\bowtie r}$ with the case where Red takes over Central Africa in the next turn. Then $\langle\langle A\rangle\rangle\, [X\psi]_{\bowtie r}$ would be true iff $\bowtie r$ is ≥ 88.98.

2. $\langle\!\langle A\rangle\!\rangle\, [\phi U\psi]_{\bowtie r}$ can represent statements such as "Red will hold all territory until North America belongs to Red".

3. $\langle\!\langle A\rangle\!\rangle\, [\phi R\psi]_{\bowtie r}$ can represent scenarios such as "Red holds Middle East until he takes over all of North and South America".

4 Model Checking

In this section, we apply model checking techniques to determine the complexity of $PATL_{ir}$. For now, we work with explicit representations of a system on this paper. Our strategy to find the complexity of $PATL_{ir}$ is similar to that of ATL_{ir}: First, we consider a linear algorithm, which sorts different cases based on the formula structure. We guess a strategy σ_A of A by an NP-Oracle. Then, the strategy is verified by pruning. We know that the result of pruning our $PATL_{ir}$ is a $PCTL$ formula as $PATL$ logics are generalisations of $PCTL$. Finally, we model-check the resulting PCTL formula.

Theorem 1 *For $\mathcal{G}$ and $PATL_{ir}$ formula φ, the $PATL_{ir}$ model-checking problem $\mathcal{M} \models \varphi$ can be solved in time Δ_2^p*

Proof: The following is a linear algorithm for the time complexity of $PATL_{ir}$ under an explicitly represented system:

Algorithm 1: Function $mcheck(\mathcal{G}, \varphi)$

Input: Finite system $\mathcal{M}$ with state set S and $PATL_{ir}$ formula φ
Output: $Sat(\varphi) = \{s \in S \mid s \models \varphi\}$
case(φ)
$\varphi \in AP$ **: return** $\lambda(p)$
$\varphi = \neg\phi$ **: return** $S \setminus mcheck(\mathcal{G}, \phi)$
$\varphi = \phi \wedge \psi$ **: return** $mcheck(\mathcal{G}, \phi) \cap mcheck(\mathcal{G}, \psi)$
$\varphi = \langle\langle A \rangle\rangle [X\phi]_{\bowtie r}$ **return:** $pre(\mathcal{G}, A, mcheck(\mathcal{G}, \phi))$
$\varphi = \langle\langle A \rangle\rangle [\phi U \psi]_{\bowtie r}$ **Guess** $\sigma_A \in \Sigma_A^{ir}$**. Then** $prune(\mathcal{M}, \sigma_A)$**.** $\mathcal{M}_{\sigma_A}$ **returns:** $\mathbb{P}_{\bowtie r}(\phi U \psi)$**. Apply** $mcheck_{PCTL}$**.**
$\varphi = \langle\langle A \rangle\rangle [\phi R \psi]_{\bowtie r}$ **Guess** $\sigma_A \in \Sigma_A^{ir}$**. Then** $prune(\mathcal{M}, \sigma_A)$**.** $\mathcal{M}_{\sigma_A}$ **returns:** $\mathbb{P}_{\bowtie r}(\phi Rpsi)$**. Apply** $mcheck_{PCTL}$**.**

Definition 7 *Function $pre(\mathcal{G}, A, S)$ is an auxiliary function that returns the exact set of states S' such that when the system is in a state $s \in S'$, agents in A can cooperate and enforce the next state to be in S consistent with $\bowtie_r$. Return $\{s \mid \exists j_A$ such that $\forall j_{Ag \setminus A},\ \delta(s, j_A, j_{Ag \setminus A}) \in S$ and $\delta(s, j_A, j_{Ag \setminus A})_{\bowtie r}\}$. We have $d_{\sigma_A}(a, s) = d(a, s)$ for $a \notin A$ and $d_{\sigma_A}(a, s) = d(a, s) \cap \sigma_a(s)$ for $a \in A$. $pre(\mathcal{G}, A, S)$*

Definition 8 *Function $prune(\mathcal{M}, \sigma_A)$ s a function that prunes our system $\mathcal{M}$, returning $\mathcal{M}_{\sigma_A} = \langle Ag, S, AP, \lambda, Act, d_{\sigma_A}, \delta_{\sigma_A}, PI, \sim_a \rangle$. If we have $d(a, s) = \{act_1, act_2\}$ and $\sigma_A(s) = act_1$, then $d_{\sigma_A}(a, s) = act_1$ and the function δ_{σ_A} is consistent with σ_A. For a move $j_a \in d(a, s)$, we have $\delta(s' \mid s, \langle j_1, \ldots, j_k \rangle) = \delta(s' \mid s\langle j_1, \ldots, j_k \rangle)$.*

$mcheck_{PCTL}$ calls to any established PCTL model-checker, like that of [5].

5 Reference

1. T. Chen, and J. Lu, Probabilistic Alternating-Time Temporal Logic and Model Checking Algorithm, Fourth International Conference on Fuzzy Systems and Knowledge Discovery, 2007.

2. X. Huang, K. Su, and C. Zhang, Probabilistic Alternating-Time Temporal Logic of Incomplete Information and Synchronous Perfect Recall, Twenty-Sixth AAAI Conference on Artificial Intelligence, 2012.

3. N. Bulling, J. Dix, and W Jamorga, Model Checking Logics of Strategic Ability: Complexity, Chapter 5 of Specification and Verification of Multi-agent Systems, 2010.

4. A. Bianco, and L. de Alfaro, Model Checking of Probabilistic and Nondeterministic Systems, Page 499-513 of Foundations of Software Technology and Theoretical Computer Science, 1995.

5. C. Baier, JP Kaoten, Probabilistic Systems, Chapter 10 of Principles of Model Checking, 2008.

6. W. Jamroga, and J. Dix, Model Checking Abilities of Agents: A Closer Look, Volume 42 Issue 3 of Theory of Computing Systems, 2007.

7. Risk Game Rules, Ultra Board Games,
https://www.ultraboardgames.com/risk/game-rules.php, 2019.

8. Line of Sight (LoS), Techopedia,
https://www.techopedia.com/definition/5069/line-of-sight-los, 2019.

9. N. Berry, RISK Analysis, Data Genetics,
http://datagenetics.com/blog/november22011/index.html, 2011.

10. Sperasort, Inc, Risk, Steam Games,
https://store.steampowered.com/app/227920/Risk/, 2013.

11. D. Parker, Markov Decision Process, Lecture 12 of Probabilistic Model Checking, Department of Computer Science Oxford University, 2011,
https://www.prismmodelchecker.org/lectures/pmc/12-mdps.pdf.